

Leaky regions

Tim Harris

Microsoft
Research

Introduction

Hypothesis: even when storage is managed by a garbage collector, it's usually possible to say when a given object can be destroyed

- GC means that:
 - The programmer doesn't have to figure this out
 - It can vary between apps, e.g. for library code
- Plus there are the oft-claimed GC advantages e.g.
 - Allocation is usually fast
 - Object retention leads to bloat rather than crashes

Introduction (ii)

- People are still worried about the cost and timing of GC
- From a Singularity discussion:

```
loop {  
    BeginTask(resource estimates, deadline);  
    ... do time-sensitive work ...  
    EndTask(out actual resources used, out slack time);  
    waitNextFrame();  
}
```


- Can we identify blocks of code like this BeginTask...EndTask block where objects allocated in the block are expected to be destroyed by the time the block ends?
- Can we exploit this to (i) reduce the cost of GC and (ii) control when work occurs?

Leaky regions

- Basic idea: annotate method calls saying “most objects allocated in this call will be dead when it returns”
- Safe: runtime checks for violations


Leaky regions

- Basic idea: annotate method calls saying “most objects allocated in this call will be dead when it returns”
- Safe: runtime checks for violations



Leaky regions

- Basic idea: annotate method calls saying “most objects allocated in this call will be dead when it returns”
- Safe: runtime checks for violations



Leaky regions

- Basic idea: annotate method calls saying “most objects allocated in this call will be dead when it returns”
- Safe: runtime checks for violations
- Three main costs in a naïve system:
 - Checking when an object leaks
 - GC-scope entry/exit code
 - Fragmentation between GC-scopes (e.g. one per page)
- Key assumption:
 - Objects rarely leak from GC-scopes

Leaky regions

- Basic idea: annotate method calls saying “most objects allocated in this call will be dead when it returns”
- Safe: runtime checks for violations
- Three main costs in a naïve system:
 - Checking when an object leaks
 - GC-scope entry/exit code
 - Fragmentation between GC-scopes (e.g. one per page)
- Key assumption:
 - Objects rarely leak from GC-scopes

Research context


- A lot of existing work on region-based collection (Tofte+Talpin '94)
- More restrictive form in RTSJ

Research context

- A lot of existing work on region-based collection (Tofte+Talpin '94)
- More restrictive form in RTSJ

Heap

- Push a new region



Research context


- A lot of existing work on region-based collection (Tofte+Talpin '94)
- More restrictive form in RTSJ



- Push a new region

Research context

- A lot of existing work on region-based collection (Tofte+Talpin '94)
- More restrictive form in RTSJ



- Push a new region
- Allocate space in any of the regions

Research context

- A lot of existing work on region-based collection (Tofte+Talpin '94)
- More restrictive form in RTSJ

Heap




- Push a new region
- Allocate space in any of the regions
- Pop & deallocate a region


Research context (ii)

- Escape analysis for on-stack allocation (OOPSLA '99)
 - As with region systems, must be proven safe
 - Per-method granularity
 - Risks of stack overflow
- Adaptive region-based allocation (Qian+Hendren, ISMM '02)
 - Mark allocation sites escaping / non-escaping
 - Won't tolerate objects that escape and are recaptured
 - Slightly better than per-method granularity
 - Looks at memory-related stats, not performance


Implementation




Implementation (ii)




Implementation (ii)



Implementation (ii)



Implementation (ii)



Annotations (xlist)

	Calls	Allocation footprint	
impact.li_130.xlist.clar-> impact.li_130.xlist.xload	20	2061013320	99%
impact.li_130.xlist.xload-> impact.li_130.xlist.xlevel	281	205878532	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.evform	41048456	205878532	99%
impact.li_130.xlist.evform-> impact.li_130.xlist.func	28756528	2058782718	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xlevel	32000688	2058710020	99%
impact.li_130.xlist.evfun-> impact.li_130.xlist.xlevel	14448211	2058704812	99%
impact.li_130.xlist.evform-> impact.li_130.xlist.evfun	2445802	2058703180	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xlevel	34218881	2055788568	99%
impact.li_130.w_xcond.func-> impact.li_130.xlist.xlevel	2703585	1994337804	96%
impact.li_130.xlist.evform-> impact.li_130.xlist.xlevel	17848325	1521302440	73%
impact.li_130.xlist.evform-> impact.li_130.xlist.xlevel	2445802	1510122408	73%
impact.li_130.xlist.tagblock-> impact.li_130.xlist.xlevel	4603864	1339567852	64%
impact.li_130.w_xdo.func-> impact.li_130.xlist.doloop	112743	1335780852	64%
impact.li_130.xlist.doloop-> impact.li_130.xlist.tagblock	2562724	1334083528	64%
impact.li_130.xlist.xgetvalue-> impact.li_130.xlist.xgetvalue	78751788	1280034382	61%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xgetvalue	78731784	1259714328	61%
impact.li_130.w_xlet.func-> impact.li_130.xlist.let	90581	1205580718	58%
impact.li_130.xlist.let-> impact.li_130.xlist.xlevel	158334	1204785380	58%
impact.li_130.w_xcond.func-> impact.li_130.xlist.xlevel	5094788	1008555580	48%
impact.li_130.w_xor.func-> impact.li_130.xlist.xlevel	54783	770584288	37%
impact.li_130.w_xprogn.func-> impact.li_130.xlist.xlevel	382864	710183472	34%

Annotations (xlist)

	Calls	Allocation footprint	
impact.li_130.xlisp.ctor -> impact.li_130.xlisp.xload		20	2061013320 99%
impact.li_130.xlisp.xload -> impact.li_130.xlisp.xlevel		281	2058795532 99%
impact.li_130.xlisp.xlevel -> impact.li_130.xlisp.evform	41048456	2058795532	99%
impact.li_130.xlisp.evform -> impact.li_130.xlisp.func	28756526	2058782718	99%
impact.li_130.xlisp.xlevel -> impact.li_130.xlisp.xlevel	32000688	2058710020	99%
impact.li_130.xlisp.evfun -> impact.li_130.xlisp.xlevel	14448211	2058704812	99%
impact.li_130.xlisp.evform -> impact.li_130.xlisp.evfun	2445802	2058703180	99%
impact.li_130.xlisp.xlevel -> impact.li_130.xlisp.xlevel	34218891	2055798568	99%
impact.li_130.w_xcond.func -> impact.li_130.xlisp.xlevel	2703585	1994337804	96%
impact.li_130.xlisp.evform -> impact.li_130.xlisp.xlevel	17848325	1521302440	73%
impact.li_130.xlisp.evform -> impact.li_130.xlisp.xlevel	2445802	1510122408	73%
impact.li_130.xlisp.tagblock -> impact.li_130.xlisp.xlevel	4603864	1339587852	64%
impact.li_130.w_xdo.func -> impact.li_130.xlisp.doloop	112743	1335780852	64%
impact.li_130.xlisp.doloop -> impact.li_130.xlisp.tagblock	2562724	1334088328	64%
impact.li_130.xlisp.xigetvalue -> impact.li_130.xlisp.xigetvalue	78751788	1280034382	61%
impact.li_130.xlisp.xlevel -> impact.li_130.xlisp.xigetvalue	78731784	1259714328	61%
impact.li_130.w_xlet.func -> impact.li_130.xlisp.let	90581	1205580716	58%
impact.li_130.xlisp.let -> impact.li_130.xlisp.xlevel	158334	1204785380	58%
impact.li_130.w_xcond.func -> impact.li_130.xlisp.xlevel	5094788	1008555580	48%
impact.li_130.w_xor.func -> impact.li_130.xlisp.xlevel	54783	770584288	37%
impact.li_130.w_xprogn.func -> impact.li_130.xlisp.xlevel	382864	710183472	34%

Annotations (xlist)

	Calls	Allocation footprint	
impact.li_130.xlist.dor-> impact.li_130.xlist.xload	20	2061013320	99%
impact.li_130.xlist.xload-> impact.li_130.xlist.xlevel	281	2058795532	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.evform	41048456	2058795532	99%
impact.li_130.xlist.evform-> impact.li_130.xlist.xlevel	20700328	2058795532	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xlevel	32000688	2058710020	99%
impact.li_130.xlist.evfun-> impact.li_130.xlist.xlevel	14448211	2058704812	99%
impact.li_130.xlist.evform-> impact.li_130.xlist.evfun	2445802	2058703180	99%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xlevel	34218881	2055798568	99%
impact.li_130.w_xcond.func-> impact.li_130.xlist.xlevel	2703585	1994337604	96%
impact.li_130.xlist.evform-> impact.li_130.xlist.xlevel	17948325	1521302440	73%
impact.li_130.xlist.evform-> impact.li_130.xlist.xlevel	2445802	1510122408	73%
impact.li_130.xlist.tagblock-> impact.li_130.xlist.xlevel	4603864	1339587852	64%
impact.li_130.w_xdo.func-> impact.li_130.xlist.doloop	112743	1335780852	64%
impact.li_130.xlist.doloop-> impact.li_130.xlist.tagblock	2562724	1334083528	64%
impact.li_130.xlist.xigetvalue-> impact.li_130.xlist.xigetvalue	78751788	1280034382	61%
impact.li_130.xlist.xlevel-> impact.li_130.xlist.xigetvalue	78731784	1259714328	61%
impact.li_130.w_xlet.func-> impact.li_130.xlist.let	90581	1205580716	58%
impact.li_130.xlist.let-> impact.li_130.xlist.xlevel	158334	1204765380	58%
impact.li_130.w_xcond.func-> impact.li_130.xlist.xlevel	5094788	1008555580	48%
impact.li_130.w_xor.func-> impact.li_130.xlist.xlevel	54783	770584288	37%
impact.li_130.w_xprogn.func-> impact.li_130.xlist.xlevel	382864	710183472	34%

Annotations (xlist)

	Calls	Allocation footprint	
impact.li_130.xlist.ctor -> impact.li_130.xlist.xload	20	2061013320	99%
impact.li_130.xlist.xload -> impact.li_130.xlist.xlevel	281	205878532	99%
impact.li_130.xlist.xlevel -> impact.li_130.xlist.evform	41048456	205878532	99%
impact.li_130.xlist.evform -> impact.li_130.xlist.func	28756528	2058782718	99%
impact.li_130.xlist.xlevel -> impact.li_130.xlist.xlevel	32000688	2058710020	99%
impact.li_130.xlist.evfun -> impact.li_130.xlist.xlevel	14449211	2058704812	99%
impact.li_130.xlist.evform -> impact.li_130.xlist.evfun	2445902	2058703180	99%
impact.li_130.xlist.xlevel -> impact.li_130.xlist.xlevel	34218881	2055788568	99%
impact.li_130.w_xcond.func -> impact.li_130.xlist.xlevel	2703585	1994337884	96%
impact.li_130.xlist.evform -> impact.li_130.xlist.xlevel	17848325	1521302440	73%
impact.li_130.xlist.evform -> impact.li_130.xlist.xlevel	2445902	1510122408	73%
impact.li_130.xlist.tagblock -> impact.li_130.xlist.xlevel	4603884	1339587852	64%
impact.li_130.w_xdo.func -> impact.li_130.xlist.doloop	112743	1335780652	64%
impact.li_130.xlist.doloop -> impact.li_130.xlist.tagblock	2562724	1334083528	64%
impact.li_130.xlist.xigetvalue -> impact.li_130.xlist.xigetvalue	78751788	1280034382	61%
impact.li_130.xlist.xlevel -> impact.li_130.xlist.xigetvalue	78731784	1259714328	61%
impact.li_130.w_xlet.func -> impact.li_130.xlist.let	90581	1205580718	58%
impact.li_130.xlist.let -> impact.li_130.xlist.xlevel	158334	1204785380	58%
impact.li_130.w_xcond.func -> impact.li_130.xlist.xlevel	5034788	1006555580	48%
impact.li_130.w_xor.func -> impact.li_130.xlist.xlevel	54783	770594288	37%
impact.li_130.w_xprogn.func -> impact.li_130.xlist.xlevel	382864	710183472	34%

Annotations (xlisp)

- `impact.li_130.w_xdo.func` -> `impact.li_130.xlisp.doloop`:
 - 112743 invocations
 - 1215MB allocated
 - 1210MB reclaimed
 - 0.4MB leaked
- `impact.li_130.xlisp.evform` -> `impact.li_130.xlisp.evfun`
 - 2445902 invocations
 - 1209MB allocated
 - 807MB reclaimed
 - 0.3MB leaked
- Many objects leaked temporarily through static fields
- Leaked objects are lisp heap segments


Annotations (crafty)

- Crafty.Net.OO.MoveGenerator.Search -> Crafty.Net.OO.MoveGenerator.ABSearch
 - 84767 invocations
 - 217MB allocated
 - 217MB reclaimed
 - 0.5MB leaked
- 87% allocation footprint
- Leaks come from checking for board-state repetitions


Annotations (go, jpeg)

- `impact.go.go.Main` -> `impact.go.go.getmove`
 - 296 invocations
 - 709MB allocation footprint
 - 704MB reclaimed
 - No leaks
 - Fairly large allocation volumes per invocation => little reclamation with small young generation
- Two calls to `go_execute_compression_and_decompression`
 - 128 invocations
 - 28MB allocation footprint
 - 28MB reclaimed
 - 0.07MB leaks (arrays in `spec_define_subimage_int`)


Results (xlisp)




Results (crafty)




Results (crafty)



Results (crafty)



Results (crafty)















Crafty

90

jpeg

xlii



Conclusions

- Practical run-time costs unlike previous work
- Evaluation shows the value of
 - Recapturing escapees
 - Considering more than individual method calls
- Picking the boundary is important
 - Few calls, little leakage, large portion of total allocation
 - Scope for feedback-directed choice
- Evaluation with other collectors
 - Other kinds of write barrier
- This is all single threaded...